CLAIMS

What is claimed is:

1 A computer implemented data consistency maintenance method comprising the steps of:

attempting to acquire a genlock on a mirror page during a write access request operation;

locking said genlock on said mirror page;

updating data on said mirror page;

syncing said mirror page to backing store;

acquiring a genlock on a primary page associated with said mirror page;

locking said genlock on said primary page;

performing an update of data on said primary page

syncing said primary page to backing store;

preventing read operations and write operations of other processes from accessing said mirror page and said primary page while locked;

unlocking said genlock; and

re-syncing said mirror page and said primary page.

2 A data consistency maintenance method of Claim 1 further comprising the steps of keeping data included in said mirror page and data included in said associated primary page in lockstep except during a write operation.

3 A data consistency maintenance method of Claim 1 further comprising the step of granting an attempt to acquire a genlock if said genlock is unlocked.

4 A data consistency maintenance method of Claim 1 further comprising the steps of:

retrieving a mirror page from a database; and

mapping said mirror page into a local memory if not already stored in said local memory.

5 A data consistency maintenance method of Claim 1 further comprising the step of:

utilizing the lowest order bit of an write counter value to function as a genlock; and

incrementing said write counter each time a write access is performed.

6 A data consistency maintenance method of Claim 1 wherein locking said genlock on said mirror page and said genlock on said primary page provides an indication that said mirror page and said primary page are being accessed by a process performing a write operation.

7 A computer implemented consistency recovery method for recovering consistency after a process crash, comprising the steps of:

comparing a write counter value of a primary page to a write counter value of an associated mirror page;

determining whether said primary page or said associated mirror page includes valid data;

copying a consistent page to an inconsistent page; and

resolving genlock status.

8 A computer implemented consistency recovery method of Claim 7 further comprising the step of allowing a writer to either complete or roll back if a process or system crash occurs while write operations are performed.

9 A computer implemented consistency recovery method of Claim 7 wherein both said primary page and said associated mirror page are consistent if said write counter value of said primary page is equal to said write counter value of said associated mirror page and said primary page and said associated mirror page are both unlocked.

10 A computer implemented consistency recovery method of Claim 7 wherein said primary page is consistent if an write counter value of said primary page is less than an write counter value of said mirror page and said primary page is unlocked and said mirror page is locked.

11 A computer implemented consistency recovery method of Claim 7 wherein said mirror page is consistent if said write counter value of said primary page is equal to said write counter value of said associated mirror page and said primary page and said associated mirror page are locked.

12 A computer implemented consistency recovery method of Claim 7 wherein said associated mirror page is consistent if said write counter value of said primary page is greater than said write counter value of said mirror page and said primary page is unlocked and said associated mirror page is locked.

13 A computer implemented consistency recovery method of Claim 7 wherein said primary page and said associated mirror page are considered to be in an invalid state if said primary page is locked and said associated mirror page is unlocked.

14 A computer implemented consistency recovery method of Claim 13 wherein data included in said associated mirror page is considered the valid information.

15 A data consistency maintenance and recovery computer system comprising:

a bus for providing a communication path between components of computer system;

a central processing unit (CPU) coupled to said bus, said CPU handles information control and data processing including a first process and a second process that perform a write operation and read operation except when prevented by a lock; and

a database coupled to said bus, said database stores data on a computer readable medium, said data arranged on a primary page and maintained in a associate mirror page that is a copy comprising said data included in said primary page except when one of said first process and said second process perform a write operation, wherein said associated mirror page is written to before said primary page during said write operation.

16 The data consistency maintenance and recovery computer system of claim 15 wherein said primary page includes a first write counter value and said mirror page includes a second write counter value.

17 The data consistency maintenance and recovery computer system of claim 16 wherein said data consistency maintenance and recovery computer system compares the status of said lock on said primary page to the status of said lock of associated said mirror page and compares said first writecounter value of said primary page to said second write counter value of said mirror

page to determine whether said primary page or said mirror page are consistent after recovering from a process or system crash.

18 The data consistency maintenance and recovery computer system of claim 15 wherein said lock is the lowest order bit of an write counter value.

19 The data consistency maintenance and recovery computer system of claim 18 wherein said lock is locked and unlocked by incrementing said write counter value.

20 The data consistency maintenance and recovery computer system of claim 15 wherein said first process and said second process lock said associated mirror page and said primary page when beginning a write operation and unlock said associated mirror page and said primary page when finishing a write operation.